

UNITED STATES PATENT APPLICATION

for

A METHOD FOR SECURING MEMORY MAPPED CONTROL REGISTERS

Inventors:

Richard P. Mangold
Keith Shippy

prepared by:

BLAKELY, SOKOLOFF, TAYLOR & ZAFMAN LLP
12400 Wilshire Boulevard
Los Angeles, CA 90025-1026
(303) 740-1980

File No.: 042390.P13346

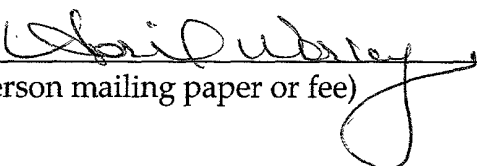
"Express Mail" mailing label number EV069174301US

Date of Deposit February 20, 2002

I hereby certify that this paper or fee is being deposited with the United States Postal Service "Express Mail Post Office to Addressee" service under 37 CFR 1.10 on the date indicated above and is addressed to the Commissioner of Patents and Trademarks, Washington, D.C. 20231.

April Worley

(Typed or printed name of person mailing paper or fee)


(Signature of person mailing paper or fee)

A METHOD FOR SECURING MEMORY MAPPED CONTROL REGISTERS

COPYRIGHT NOTICE

[0001] Contained herein is material that is subject to copyright protection. The copyright owner has no objection to the facsimile reproduction of the patent disclosure by any person as it appears in the Patent and Trademark Office patent files or records, but otherwise reserves all rights to the copyright whatsoever.

FIELD OF THE INVENTION

[0002] The present invention relates to computer systems; more particularly, the present invention relates to the protection of data received at a computer system.

BACKGROUND

[0003] Widespread use of the Internet and electronic mail (e-mail) has left millions of personal computers (PCs) vulnerable to downloaded viruses and other types of malicious software that can destroy programs, copy and upload private documents, and perform other harmful acts, frequently without the PC operator's knowledge. The increasing popularity of downloaded programs has multiplied the problem significantly, since such programs create more opportunities to unknowingly download the malicious software.

[0004] Due to their open architecture, most PCs provide very little protection against such destructive software. It is this very openness that has

made the PC platform the general-purpose solution provider that it is. Other types of computers are also vulnerable to such attacks in varying degrees, but the pervasive use of PCs has drawn much attention to the problem as it applies to PCs.

[0005] In the past, owners of copyrighted information or other intellectual property have been reluctant to allow their property to be viewed on the PC platform (books, movies, sensitive corporate documents, etc.) as the nature of the open PC platform makes the property vulnerable to mischievous software that may run in the background. Although self-replicating destructive software (viruses) attracts the most attention, copyright owners are more concerned with the illegal copying and distribution of any document that they permit to be downloaded to a computer. This is particularly true of video that is available electronically by downloading over a network such as the Internet, or service provider networks. The ease of copying downloaded video makes it easy to illicitly reproduce and forward copyrighted materials without detection of this activity by the copyright owner.

BRIEF DESCRIPTION OF THE DRAWINGS

[0006] The present invention will be understood more fully from the detailed description given below and from the accompanying drawings of various embodiments of the invention. The drawings, however, should not be taken to limit the invention to the specific embodiments, but are for explanation and understanding only.

[0007] Figure 1 illustrates one embodiment of a network;

[0008] Figure 2 is a block diagram of one embodiment of a computer system;

[0009] Figure 3 is a block diagram of one embodiment of a mechanism for protecting memory mapped control registers from unauthorized access; and

[0010] Figure 4 is a flow diagram for one embodiment of protecting memory mapped control registers from unauthorized access.

DETAILED DESCRIPTION

[0011] A method for protecting memory mapped control registers from unauthorized access by user level applications is described. According to one embodiment, a device driver is initialized upon powering up a computer system or set-top appliance. During initialization, the driver locates a video decoder and allocates any detected physical hardware registers. Subsequently, the driver passes the list of hardware registers to a kernel security module.

[0012] The kernel security module verifies the digital signature of the driver. The kernel security module then securely stores the list for later use. Afterward, an application program begins to receive video content from a content source. The application calls an application level interface to stream content to the video decoder. It is then determined whether the application call is the first call for video streaming.

[0013] If it is determined that it is the first call, the interface requests the list of hardware registers from the device driver. The driver then maps the resources into a virtual space accessible from the application and returns the resources. The application level interface forwards the virtual memory address to an application level security module to validate the hardware registers to which the application intends to stream the content.

[0014] Subsequently, the application security module verifies the signature of the interface and calls the kernel security module via a secure interface in order to verify the application's virtual resources map. As a result,

the application's virtual resources map is compared to the stored list of physical resources. If the resources match, the video content is streamed through those resources directly to the video decoder for processing. Otherwise, the video content will be precluded from steaming to the video decoder.

[0015] Reference in the specification to "one embodiment" or "an embodiment" means that a particular feature, structure, or characteristic described in connection with the embodiment is included in at least one embodiment of the invention. The appearances of the phrase "in one embodiment" in various places in the specification are not necessarily all referring to the same embodiment.

[0016] In the following description, numerous details are set forth. It will be apparent, however, to one skilled in the art, that the present invention may be practiced without these specific details. In other instances, well-known structures and devices are shown in block diagram form, rather than in detail, in order to avoid obscuring the present invention.

[0017] Some portions of the detailed descriptions that follow are presented in terms of algorithms and symbolic representations of operations on data bits within a computer memory. These algorithmic descriptions and representations are the means used by those skilled in the data processing arts to most effectively convey the substance of their work to others skilled in the art. An algorithm is here, and generally, conceived to be a self-consistent sequence of steps leading to a desired result. The steps are those requiring physical manipulations of

physical quantities. Usually, though not necessarily, these quantities take the form of electrical or magnetic signals capable of being stored, transferred, combined, compared, and otherwise manipulated. It has proven convenient at times, principally for reasons of common usage, to refer to these signals as bits, values, elements, symbols, characters, terms, numbers, or the like.

[0018] It should be borne in mind, however, that all of these and similar terms are to be associated with the appropriate physical quantities and are merely convenient labels applied to these quantities. Unless specifically stated otherwise as apparent from the following discussion, it is appreciated that throughout the description, discussions utilizing terms such as "processing" or "computing" or "calculating" or "determining" or "displaying" or the like, refer to the action and processes of a computer system, or similar electronic computing device, that manipulates and transforms data represented as physical (electronic) quantities within the computer system's registers and memories into other data similarly represented as physical quantities within the computer system memories or registers or other such information storage, transmission or display devices.

[0019] The present invention also relates to an apparatus for performing the operations herein. This apparatus may be specially constructed for the required purposes, or it may comprise a general-purpose computer selectively activated or reconfigured by a computer program stored in the computer. Such a computer program may be stored in a computer readable storage medium, such

as, but is not limited to, any type of disk including floppy disks, optical disks, CD-ROMs, and magnetic-optical disks, read-only memories (ROMs), random access memories (RAMs), EPROMs, EEPROMs, magnetic or optical cards, or any type of media suitable for storing electronic instructions, and each coupled to a computer system bus.

[0020] The algorithms and displays presented herein are not inherently related to any particular computer or other apparatus. Various general-purpose systems may be used with programs in accordance with the teachings herein, or it may prove convenient to construct more specialized apparatus to perform the required method steps. The required structure for a variety of these systems will appear from the description below. In addition, the present invention is not described with reference to any particular programming language. It will be appreciated that a variety of programming languages may be used to implement the teachings of the invention as described herein.

[0021] The instructions of the programming language(s) may be executed by one or more processing devices (e.g., processors, controllers, control processing units (CPUs), execution cores, etc.).

[0022] **Figure 1** illustrates one embodiment of a network 100. Network 100 includes a computer system 110 and a computer system 120 coupled via a transmission medium 130. In one embodiment, computer system 110 operates as a source device that sends an object to computer system 120, operating as a receiving device. The object may be, for example, a video file, a data file, an

executable, or other digital objects. The object is sent via data transmission medium 130.

[0023] The data transmission medium 130 may be one of many mediums such as a satellite transmission, an internal network connection, an Internet connection, or other connections. The transmission medium 130 may be connected to a plurality of untrusted routers (not shown) and switches (not shown).

[0024] According to one embodiment, computer system 110 is a server computer, while computer system 120 is a client set-top appliance. In a further embodiment, the set-top appliance is implemented for cable television or digital satellite services that receive audio and video content from computer system 110. However, in other embodiments, computer system 120 may be a personal computer (PC) coupled to computer system 110, wherein the transmission medium is the Internet.

[0025] Figure 2 is a block diagram of one embodiment of a computer system 200. Computer system 200 may be implemented as computer system 120 (both shown in Figure 1). The computer system 200 includes a processor 201 that processes data signals. Processor 201 may be a complex instruction set computer (CISC) microprocessor, a reduced instruction set computing (RISC) microprocessor, a very long instruction word (VLIW) microprocessor, a processor implementing a combination of instruction sets, or other processor device.

200220-10067001

[0026] In one embodiment, processor 201 is a processor in the Pentium® family of processors including the Pentium® II family and mobile Pentium® and Pentium® II processors available from Intel Corporation of Santa Clara, California. Alternatively, other processors may be used. **Figure 2** shows an example of a computer system 200 employing a single processor computer. However, one of ordinary skill in the art will appreciate that computer system 200 may be implemented using multiple processors.

[0027] Processor 201 is coupled to a processor bus 210. Processor bus 210 transmits data signals between processor 201 and other components in computer system 200. Computer system 200 also includes a memory 213. In one embodiment, memory 213 is a dynamic random access memory (DRAM) device. However, in other embodiments, memory 213 may be a static random access memory (SRAM) device, or other memory device.

[0028] Memory 213 may store instructions and code represented by data signals that may be executed by processor 201. According to one embodiment, a cache memory 202 resides within processor 201 and stores data signals that are also stored in memory 213. Cache 202 speeds up memory accesses by processor 201 by taking advantage of its locality of access. In another embodiment, cache 202 resides external to processor 201.

[0029] Computer system 200 further comprises a bridge memory controller 211 coupled to processor bus 210 and memory 213. Bridge/memory controller 211 directs data signals between processor 201, memory 213, and other

components in computer system 200 and bridges the data signals between processor bus 210, memory 213, and a first input/output (I/O) bus 220.

[0030] In one embodiment, I/O bus 220 may be a single bus or a combination of multiple buses. In a further embodiment, I/O bus 220 may be a Peripheral Component Interconnect adhering to a Specification Revision 2.1 bus developed by the PCI Special Interest Group of Portland, Oregon. In another embodiment, I/O bus 220 may be a Personal Computer Memory Card International Association (PCMCIA) bus developed by the PCMCIA of San Jose, California. Alternatively, other busses may be used to implement I/O bus. I/O bus 220 provides communication links between components in computer system 200.

[0031] A network controller 221 is coupled I/O bus 220. Network controller 221 links computer system 200 to a network of computers (not shown in **Figure 2**) and supports communication among the machines. In one embodiment, computer system 200 receives streaming video data from a computer 110 via network controller 221.

[0032] A display device controller 222 is also coupled to I/O bus 220. Display device controller 222 allows coupling of a display device to computer system 200, and acts as an interface between the display device and computer system 200. In one embodiment, display device controller 222 is a monochrome display adapter (MDA) card.

[0033] In other embodiments, display device controller 222 may be a color

graphics adapter (CGA) card, an enhanced graphics adapter (EGA) card, an extended graphics array (XGA) card or other display device controller. The display device may be a television set, a computer monitor, a flat panel display or other display device. The display device receives data signals from processor 201 through display device controller 222 and displays the information and data signals to the user of computer system 200.

[0034] A video decoder 223 is also coupled to I/O bus 220. Video decoder 223 is a hardware device that translates received encoded data into its original format. According to one embodiment, video decoder 223 is a Moving Picture Expert Group 4 (MPEG-4) decoder. However, one of ordinary skill in the art will appreciate that video decoder 223 may be implemented with other types of MPEG decoders.

[0035] Computer system 200 also includes a second I/O bus 230 coupled to I/O bus 220 via a bus bridge 224. Bus bridge 224 operates to buffer and bridge data signals between I/O bus 220 and I/O bus 230. I/O bus 230 may be a single bus or a combination of multiple buses. In one embodiment, I/O bus 230 is an Industry Standard Architecture (ISA) Specification Revision 1.0a bus developed by International Business Machines of Armonk, New York. However, other bus standards may also be used, for example Extended Industry Standard Architecture (EISA) Specification Revision 3.12 developed by Compaq Computer, et al.

[0036] I/O bus 230 provides communication links between components in

computer system 200. A data storage device 231 is coupled to I/O bus 230. I/O device 231 may be a hard disk drive, a floppy disk drive, a CD-ROM device, a flash memory device or other mass storage device. A keyboard interface 232 is also coupled to I/O bus 230. Keyboard interface 232 may be a keyboard controller or other keyboard interface. In addition, keyboard interface 232 may be a dedicated device or can reside in another device such as a bus controller or other controller. Keyboard interface 232 allows coupling of a keyboard to computer system 200 and transmits data signals from the keyboard to computer system 200. An audio controller is also coupled to I/O bus 230. Audio controller 233 operates to coordinate the recording and playing of sounds.

[0037] According to one embodiment, computer system 200 includes a mechanism that secures received video data from unauthorized access. In many operating systems, software applications are provided access to hardware control registers as mapped into the particular application's memory space. In particular, an application may be granted access to video decoder registers for a video decoder such as video decoder 223. In such an instance, an MPEG source may stream content into the video decoder hardware by accessing the video decoder registers using libraries provided by a hardware vendor.

[0038] Typically, there is a window of opportunity for a rogue application to access the video content between the time the video decoder is queued to access the memory registers and the time the content is decoded and executed by the video decoder. During that window, a rogue application may give a false

memory address to read or write to (e.g., intercept the call to the video decoder and trick the application software into reading or writing the content to memory registers associated with another software application).

[0039] In addition, using virtual memory technology, the memory address may remain the same. However, the page tables may be changed such that the memory addresses are correct, but physically translate to a different address. Therefore, a mechanism and apparatus to protect against an un-authorized application deceiving an MPEG source application into writing the content into registers that are not actually located in video decoder hardware is provided.

[0040] **Figure 3** is a block diagram of one embodiment of a mechanism to protect memory mapped control registers from unauthorized access. According to one embodiment, the mechanism is implemented at the application space and the kernel space. The application space components include a content source 310, a software interface 320 and an application security module 350. The kernel space components include a video decoder device driver 330, video decoder 223 and a kernel security module 340.

[0041] Content source 310 transmits digital video content to interface 320. As discussed above, content source 310 may be a public network (e.g., the Internet), or a private network. In other embodiments, content source 310 may be a DVD or other video source. Software interface 320 serves as an interface to a software player application (not shown) that plays the video content received

from content source 310. In particular, interface 320 accesses and decrypts the received video content.

[0042] Device driver 330 is a software program that interacts directly with video decoder 223. In particular, driver 330 enumerates and communicates with decoder 223. In one embodiment, driver 330 detects and allocates hardware registers within video decoder 223 for access by applications at computer system 200 boot up. In addition, driver 330 passes a list of any allocated resources to tamper resistant module 340. According to a further embodiment, driver 330 and interface 320 implement digital signatures and public/private key technologies to prevent tampering and to ensure that the driver 330/decoder 223 combination have authorized access to the video content.

[0043] Application tamper resistant security module 340 receives the list of allocated resources from driver 330. Upon receiving the resources, security module 340 verifies the integrity of driver 330 via digital signatures. In addition, security module 340 securely stores the list of allocated resources.

[0044] Application tamper resistant security module 350 receives a list of allocated application level resources from interface 320 whenever a streaming application (or a library that interfaces to the hardware at the application level) is about to release protected content to registers within video decoder 223. In response, security module 350 will make a call to security module 340 which will translate the virtual (application) level resources to physical resources and compare the resources against a stored list. According to one embodiment, the

interface between security module 340 and security module 350 is secured by a random number secret key system to prevent an unauthorized application from deceiving security module 350 into relinquishing the list of resources.

[0045] Figure 4 is a flow diagram for one embodiment of protecting memory mapped control registers from unauthorized access. At processing block 405, device driver 330 is initialized upon computer system 200 boot up. In particular, driver 330 locates video decoder 223 and allocates any detected physical resources (e.g., physical hardware registers).

[0046] At processing block 410, driver 330 passes the list of resources to kernel security module 340. At processing block 415, security module 340 verifies the digital signature of the calling driver 330. At processing block 420, security module 340 securely stores the list for later use. As a result, security module 340 encrypts the list and prevents unauthorized access.

[0047] At processing block 425, an application program, not shown, begins to receive video content from content source 310. At processing block 430, the application calls application level interface 320 to stream content to video decoder 223. At decision block 435, it is determined the application call is the beginning of streaming (first call).

[0048] If it is determined that it is the beginning of streaming, interface 320 requests resources from device driver 330, processing block 440. At processing block 445, driver 330 maps the resources into a virtual space accessible from the application, and returns the resources. At processing block 450, application level

interface 320 calls (e.g., forwards the virtual memory address) application level security module 350 to validate the resources to which the application intends to stream the content.

[0049] At processing block 455, application security module 350 verifies the signature of interface 320. At processing block 460, application security module 350 calls kernel security module 340 in order to verify the application's virtual resources map. As a result, the application's virtual resources map is compared to the stored list of physical resources.

[0050] At decision block 465, it is determined whether the resources match. If the resources match, the video content is streamed through those resources directly to video decoder 223 for processing, processing block 470. Otherwise, the video content will be precluded from steaming to video decoder 223.

[0051] If at decision block 435, it is determined that it is not the beginning of streaming, control is forwarded to processing block 450 where application level interface 320 calls application level security interface 350 to validate the resources to which it intends to stream the content.

[0052] The above-described method and mechanism enables a secure video playback application to ensure that it is streaming content through resources that were allocated by a known legitimate kernel level device driver. The device driver will allocate and secure these resources at boot time, making it difficult to modify the resource list at run time.

[0053] Whereas many alterations and modifications of the present invention will no doubt become apparent to a person of ordinary skill in the art after having read the foregoing description, it is to be understood that any particular embodiment shown and described by way of illustration is in no way intended to be considered limiting. Therefore, references to details of various embodiments are not intended to limit the scope of the claims which in themselves recite only those features regarded as the invention.

2002220-40062001